

# **A Prototype for Ground-based Automated Rover Command Generation**

**Rob Sherwood, Tara Estlin, Steve Chien, Gregg Rabideau,**

**Barbara Engelhardt, Andrew Mishkin, Brian Cooper**

Jet Propulsion Laboratory  
4800 Oak Grove Dr.  
Pasadena, CA 91109  
818-393-5378  
rob.sherwood@jpl.nasa.gov

## **Abstract**

This paper will discuss a proof-of-concept prototype for automatic generation of validated rover command sequences from high-level science and engineering activities. This prototype is based on ASPEN, the Automated Scheduling and Planning Environment. This AI-based planning and scheduling system will automatically generate a command sequence that will execute within resource constraints and satisfy flight rules. Commanding the rover to achieve mission goals requires significant knowledge of the rover design, access to the low-level rover command set, and an understanding of the performance metrics rating the desirability of alternative sequences. It also requires coordination with external events such as orbiter passes and day/night cycles. An automated planning and scheduling system encodes this knowledge and uses search and reasoning techniques to automatically generate low-level command sequences while respecting rover operability constraints, science and engineering preferences, and also adhering to hard temporal constraints. Enabling goal-driven commanding of planetary rovers by engineering and science personnel greatly reduces the requirements for highly skilled rover engineering personnel and Rover Science Team time. This in turn greatly reduces mission operations costs. In addition, goal-driven commanding permits a faster response to changes in rover state (e.g., faults) or science discoveries by removing the time consuming manual sequence validation process, allowing rapid what-if analyses, and thus reducing overall cycle times.

## **Introduction**


Unlike more traditional deep space missions, surface roving missions must be operated in a reactive mode, with mission planners waiting for an end of day telemetry downlink--including critical image data--in order to plan

the next day's worth of activities. Communications time delays over interplanetary distances preclude simple 'joysticking' of the rover. A consequence of this approach to operations is that the full cycle of telemetry receipt, science and engineering analysis, science plan generation, command sequence generation and validation, and uplink of sequence, must typically be performed in twelve hours or less. Yet current rover sequence generation is manual, with limited ability to automatically generate valid rover activity sequences from more general activities/goals input by science and engineering team members. Tools such as the Rover Control Workstation (RCW) and the Web Interface for Telescience (WITS) provide mechanisms for human operators to manually generate plans and command sequences. These tools even estimate some types of resource usage and identify certain flight rule violations. However, they do not provide any means to modify the plan in response to the constraints imposed by available resources or flight rules, except by continued manual editing of sequences. This current situation has two drawbacks. First, the operator-intensive construction and validation of sequences puts a tremendous workload on the rover engineering team. The manual process is error-prone, and can lead to operator fatigue over the many months of mission operations. Second, the hours that must be reserved for sequence generation and validation reduces the time available to the science team to identify science targets and formulate a plan for submission to the engineering team. This results in reduced science return. An automated planning tool would allow the science team and sequence team to work together to optimize the plan. Many different plan options can be explored. The faster turnaround of automated planning also permits shorter than once a day planning cycles.

The Rover Control Workstation (RCW) tool, used to operate the Sojourner rover during the Pathfinder mission, does provide visualization for vehicle traverse planning, a command interface, constraint checking for individual

commands, and some resource estimation (for sequence execution time and telemetry volume). However, this tool was never intended for automated goal-based planning of rover activities. To deal with these issues, there is a need for a new tool for automated goal-based planning of rover activities.

We are using AI planning/scheduling technology to automatically generate valid rover command sequences from activity sequences specified by the mission science and engineering team. This system will automatically generate a command sequence that will execute within resource constraints and satisfy flight rules. Commanding the rover to achieve mission goals requires significant knowledge of the rover design, access to the low-level rover command set, and an understanding of the performance metrics rating the desirability of alternative sequences. It also requires coordination with external events such as orbiter passes and day/night cycles. An automated planning and scheduling system encodes this knowledge and uses search and reasoning techniques to automatically generate low-level command sequences while respecting rover operability constraints, science and engineering preferences, and also adhering to hard temporal constraints. A ground-based interactive planner combines the power of automated reasoning and conflict resolution techniques with the insights of the PI to prioritize and re-prioritize mission goals.




**Figure 1 - Mars 2001 Lander**

### Mars Surveyor Lander

The Mars Surveyor 2001 Lander is scheduled for launch on April 10, 2001. It will land on Mars on Jan. 22, 2002, if launched on schedule. The 2001 Lander will carry an imager to take pictures of the surrounding terrain during its' rocket-assisted descent to the surface. The descent-imaging camera will provide images of the landing site for geologic analyses, and will aid planning for initial operations and traverses by the rover. The 2001 Lander will also be a platform for instruments and technology experiments designed to provide key insights to decisions regarding successful and cost-effective human missions to

Mars. Hardware on the Lander will be used for an in-situ demonstration test of rocket propellant production using gases in the Martian atmosphere. Other equipment will characterize the Martian soil properties and surface radiation environment. Figure 1 contains a diagram of the lander and instruments. The Marie Curie rover will be deployed using a robotic-arm contained on the lander.



**Figure 2 - Marie Curie Rover**

The Marie Curie rover is very similar to the Mars Pathfinder Sojourner rover. (See Figure 2.) In fact, it is the same rover that was used in the Pathfinder test bed during the mission. (Mishkin et al., 1998; Mishkin 1998) Additional modifications were made to accommodate the robotic-arm-based deployment from the 2001 Lander. In addition, some minor engineering enhancements were added. A description of the rover components is contained in Table 1.

- ◆ 6-Wheeled robotic vehicle, rocker-bogie mobility chassis
- ◆ Mass: 10.5 kilograms
- ◆ Deployed volume: 65cm (l) by 48cm (w) by 30cm (h).
- ◆ Intel 80C85 CPU (~100Kips), 16K PROM, 64K rad hard RAM, 176K EEPROM, 512K RAM
- ◆ Forward Black & White stereo cameras, and rear B&W mono camera
- ◆ GaAs solar panel (16W peak)
- ◆ Primary (non-rechargeable) batteries
- ◆ UHF Radio Modem
- ◆ Laser stripers for hazard detection

**Table 1 - Rover Description**

## ASPEN Planning System

Planning and scheduling technology offers considerable promise in automating rover operations. Planning and scheduling rover operations involves generating a sequence of low-level commands from a set of high-level science and engineering goals.

ASPEN (Fukanaga, et al., 1997; Rabideau, et al., 1999) is a re-configurable planning and scheduling software framework that includes the following set of software components:

- ◆ An expressive constraint modeling language to allow the user to define naturally the application domain
- ◆ A constraint management system for representing and maintaining spacecraft operability and resource constraints, as well as activity requirements
- ◆ A set of search strategies
- ◆ A temporal reasoning system for expressing and maintaining temporal constraints
- ◆ A graphical user interface for visualizing plans/schedules

In ASPEN, the main algorithm for automated planning and scheduling is based on a technique called *iterative repair* (Zweben et al., 1994). During iterative repair, the conflicts in the schedule are detected and addressed one at a time until no conflicts exist, or a user-defined time limit has been exceeded. A conflict is a violation of a reservation, parameter dependency or temporal constraint. Conflicts can be repaired by means of several predefined methods. The repair methods are: moving an activity, adding a new instance of an activity, deleting an activity, detailing an activity, abstracting an activity, making a reservation of an activity, canceling a reservation, connecting a temporal constraint, disconnecting a constraint, and changing a parameter value. The repair algorithm may use any of these methods in an attempt to resolve a conflict. How the algorithm works is largely dependent on the type of conflict being resolved.

Rover knowledge is encoded in ASPEN under seven core model classes: activities, parameters, parameter dependencies, temporal constraints, reservations, resources and state variables. An activity is an occurrence over a time interval that in some way affects the rover. It can represent anything from a high-level goal or request to a low-level event or command. Activities are the central structures in ASPEN, and also the most complicated. Together, these constructs can be used to define rover procedures, rules and constraints in order to allow manual or automatic generation of valid sequences of activities, also called plans or schedules.

Once the types of activities are defined, specific instances can be created from the types. Multiple activity instances created from the same type might have different parameter values, including the start time. Many camera-imaging activities, for example, can be created from the

same type but with different image targets and at different start times. The sequence of activity instances is what defines the plan.

The flight rules and constraints are defined within the activities. The flight rules can be defined as temporal constraints, resource constraints, or system state constraints. Temporal constraints are defined between activities. An example would be that the rate sensor must warm up for two to three minutes before traversing (moving) the rover. In ASPEN, this would be modeled within the "move rover" activity as shown in Figure 3. The rate sensor\_heat\_up is another activity that is presumed to turn on a rate sensor heater.

Constraints can also be state or resource related. State constraints can either require a particular state or change to a particular state. Resource constraints can use a particular amount of a resource. Resources with a capacity of one are called atomic resources. ASPEN also uses non-depletable and depletable resources. Non-depletable resources are resources that can be used by more than one activity at a time and do not need to be replenished. Each activity can use a different quantity of the resource. An example would be the rover solar array power. Depletable resources are similar to non-depletable except that their capacity is diminished after use. In some cases their capacity can be replenished (memory capacity) and in other cases it cannot (battery energy, i.e. non-rechargeable primary batteries). Resource and state constraints are defined within activities using the keyword "reservations." See Figure 3 for an example.

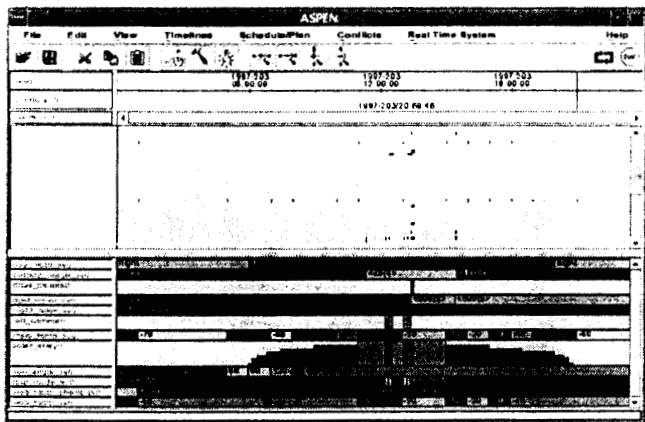
```
Activity move_rover {  
  constraints =  
    starts after end_of rate_sensor_heat_up by [2m,3m];  
  reservations =  
    solar_array_power use 35,  
    rate_sensor_state change to "on",  
    target_state must_be "ready";  
};
```

Figure 3 - ASPEN Modeling Language Example

The job of a planner/scheduler, whether manual or automated, is to accept high-level goals and generate a set of low-level activities that satisfy the goals and do not violate any of the rover flight rules or constraints. ASPEN provides a Graphical User Interface (GUI) for manual generation and/or manipulation of activity sequences. Figure 4 contains a screen dump of the GUI.

## Model Description

The Marie Curie model was built to a level at which all flight rules and constraints could be implemented. The resources include the three cameras, Alpha Proton X-Ray Spectrometer (APXS), APXS deploy motor, drive motors, solar array, battery, RAM usage, and EEPROM usage.



**Figure 4 - ASPEN GUI**

There are 27 different state variables used to track the status of various devices, modes, and parameters. Some of these parameters map directly onto rover internal parameters and others are related to the ASPEN-specific model. We are not modeling all rover internal parameters because many are not useful for automating planning. We have defined 162 activities of which 63 decompose directly into low-level rover commands.

There are several constraints that affect overall operations of the Marie Curie rover. These include:

- ◆ Earth-Mars one-way communications time delay (5-20 minutes)
- ◆ Limited communications bandwidth (generally < 10 Mbits downlink per sol available to rover)
- ◆ Limited communications opportunities (1 command uplink, 2 telemetry downlinks per sol)

The power system is the single most important resource for the Marie Curie Rover. This system consists of a .22 square meter solar array and 9 LiSOCL batteries. The batteries on Marie Curie are primarily used during the night for APXS data collection. They are primary batteries and therefore modeled as non-renewable depletable resources. The solar array is the primary power source used during the day. The predicted available solar power profile throughout the Mars day must be input before planning begins. Using a daily model is required due to changing solar array power available as a result of degradation from dust accumulation and seasonal solar irradiation variability.

A typical Mars day might involve a subset of the following activities:

- ◆ Complete an APXS data collection that was carried out during the prior night
- ◆ Capture a rear image of the APXS site
- ◆ Traverse to an appropriate site and perform a series of soil mechanics experiments, including several subframe images of soil mounds and depressions created by running individual wheel motors
- ◆ Traverse to a designated rock or soil location
- ◆ Place the APXS sensor head

- ◆ Capture end-of-day operations images with its forward cameras
- ◆ Begin APXS data collection
- ◆ Shut down for the night

APXS data collection usually occurs overnight while the rover is shutdown. Each of these activities can be input into ASPEN as a goal for that Mars day planning horizon. The format of the input goals is RML or Rover Modeling Language. RML is an application of Extensible Markup Language (XML) designed specifically for rover operations. RCW will use RML for input and output. Using a common language between the various operations tools simplifies the interface between tools.

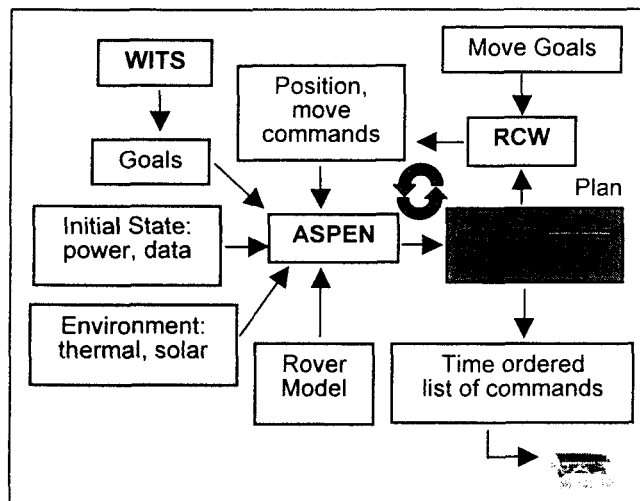
The exact position of the rover after a traverse activity is subject to dead reckoning error. The timing of traverse activities is also non-determinant. Because of the inherent problems of coordinating activities between the event-based rover and time-based lander, wait commands are used to synchronize activities. When the lander is imaging the rover after a traverse, a wait command is used to ensure the rover will remain stationary at its' destination until the lander completes imaging. Because the rover executes commands serially, this ensures that another command will not start execution before the previous command has completed. All rover traverse goals are generated using the RCW. (ASPEN is not designed to generate rover motion planning.) The RCW operator can fly a 3-D rover icon through the stereoscopic display of the Martian terrain. By inspecting the stereo scene, as well as placing the rover icon in various positions within the scene, the operator can assess the trafficability of the terrain. By placing the icon in the appropriate position and orientation directly over the stereo image of the actual rover on the surface, the rover's location and heading are automatically computed. This position information is output to ASPEN to set the rover end position state. The rover driver specifies the rover's destinations by designating a series of waypoints in the scene, generating waypoint traverse commands. The traverse commands are only a small fraction of most command sequences.

Rover data storage is a scarce resource that must be tracked within the ASPEN model. The largest consumer of data storage is the imaging data taking activity. This activity can fill the on-board data storage if a telemetry session with the lander is not available during the data collection. ASPEN will keep track of the data storage resource to ensure that all data is downlinked before the buffer is completely full.

## **Status**

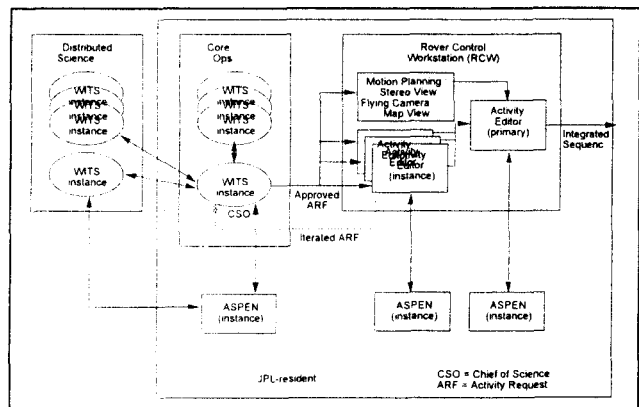
Initial work in 1998 consisted of a preliminary proof of concept demonstration in which we used automated planning and scheduling technology integrated with WITS to demonstrate automated commanding for the Rocky-7 rover from the WITS interface. In 2000, we are providing an in-depth validation of the automated command-

The end-to-end data flow for this system is shown in Figure 5. The interaction between ASPEN and RCW is an iterative process. Both ASPEN and RCW will receive high-level goals. RCW will generate initial traverse commands for input into ASPEN. ASPEN will merge these with other goals to produce an intermediate level plan. The plan will be output to RCW. This process will continue until an acceptable plan is generated. Finally a time ordered list of commands will be output for sequence generation.



The Marie Curie ASPEN model is nearly complete and ready for testing. Initial testing on a sample of 136 activities produced a conflict free plan in about 9 seconds. This testing was completed on a Sun Ultra-2 workstation. Relatively quick plan cycles will allow the Marie Curie Rover operations team to perform what-if analysis on different daily plans. In addition, the capability could be used to generate commands more frequently than once-per-day, if communications opportunities permit.

Once the model is validated, we will integrate ASPEN with RCW and WITS. Figure 6 shows a possible Marie Curie rover uplink operational data flow. The highlighted boxes show the planner that would be used at both the science planning and engineering planning level. The planner model would contain sufficient engineering information to ensure that the vast majority of science requests finally approved are feasible from an engineering standpoint. Eventually we would like to add performance metrics to the planner model to optimize the generated plans. This will enable automated "what-if" analysis to generate plans that maximize science and engineering value.



## Related Work

Once a plan has been generated it is continuously updated during plan execution to correlate with sensor and other feedback from the environment. In this way, the planner is highly responsive to unexpected changes, such as an unexpected fortuitous event or equipment failure, and can quickly modify the plan as needed. For example, if the

rover wheel slippage has caused the position estimate uncertainty to grow too large, the planner can immediately command the rover to stop and perform localization earlier than originally scheduled. Or, if a particular traversal has used more battery power than expected, the planner may need to discard one of the remaining science goals. CASPER has been integrated with control software from the JPL Rocky 7 rover (Volpe et al., 2000) and is currently being tested on Rocky 7 in the JPL Mars Yard.

## Conclusions

Current approaches to rover-sequence generation and validation are largely manual, resulting in an expensive, labor and knowledge intensive process. This is an inefficient use of scarce science-PI and key engineering staff resources. Automation as targeted by this system would automatically generate a constraint and flight rule checked time ordered list of commands and provides resource analysis options to enable users to perform more informative and fast trade-off analyses. Additionally, this technology would coordinate sequence development between science and engineering teams and would thus help speed up the consensus process.

Enabling goal-driven commanding of planetary rovers by engineering and science personnel greatly reduces the workforce requirements for highly skilled rover engineering personnel. The reduction in team size in turn greatly reduces mission operations costs. In addition, goal-driven commanding permits a faster response to changes in rover state (e.g., faults) or science discoveries by removing the time consuming manual sequence validation process, allowing rapid what-if analyses, and thus reducing overall cycle times.

## Acknowledgement

The research described in this paper was funded by the TMOD Technology Program and performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

## References

- S. Chien, R. Knight, A. Stechert, R. Sherwood, G. Rabideau, "Integrated Planning and Execution for Autonomous Spacecraft", Proceedings of the IEEE Aerospace Conference (IAC), Aspen, CO, March 1999.
- A. Fukunaga, G. Rabideau, S. Chien, D. Yan, "Toward an Application Framework for Automated Planning and Scheduling," Proceedings of the 1997 International Symposium on Artificial Intelligence, Robotics and Automation for Space, Tokyo, Japan, July 1997.
- A. Mishkin, "Field Testing on Mars: Experience Operating the Pathfinder Microrover at Ares Vallis," presentation at Field Robotics: Theory and Practice workshop, May 16 1998, at the 1998 IEEE International Conference on Robotics and Automation, Leuven, Belgium.
- A. Mishkin, J. Morrison, T. Nguyen, H. Stone, B. Cooper, B. Wilcox, "Experiences with Operations and Autonomy of the Mars Pathfinder Microrover," proceedings of the 1998 IEEE Aerospace Conference, March 21-28 1998, Snowmass at Aspen, Colorado.
- G. Rabideau, R. Knight, S. Chien, A. Fukunaga, A. Govindjee, "Iterative Repair Planning for Spacecraft Operations in the ASPEN System," International Symposium on Artificial Intelligence Robotics and Automation in Space (ISAIRAS), Noordwijk, The Netherlands, June 1999.
- R. Volpe, T. Estlin, S. Laubach, C. Olson, and J. Balaram, "Enhanced Mars Rover Navigation Techniques" To appear in the Proceedings of the IEEE International Conference on Robotics and Automation, San Francisco, CA, April 2000.
- Zweben, M., Daun, B., Davis, E., and Deale, M., "Scheduling and Rescheduling with Iterative Repair," Intelligent Scheduling, Morgan Kaufmann, San Francisco, 1994, pp. 241-256.